



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/608,037	06/30/2003	Sanjay Ghemawat	0026-0031	8255
44989 7590 10/31/2007 HARRITY SNYDER, LLP 11350 Random Hills Road SUITE 600 FAIRFAX, VA 22030			EXAMINER LOVEL, KIMBERLY M	
			ART UNIT 2167	PAPER NUMBER
			MAIL DATE 10/31/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

<b>Office Action Summary</b>	Application No. 10/608,037	Applicant(s) GHEMAWAT ET AL.	
	Examiner Kimberly Lovel	Art Unit 2167	

**-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --**

**Period for Reply**

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

**Status**

- 1) ☒ Responsive to communication(s) filed on 02 August 2007.
- 2a) ☐ This action is **FINAL**.                      2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

**Disposition of Claims**

- 4) ☒ Claim(s) 1,3-11,13,15,16 and 19-33 is/are pending in the application.
- 4a) Of the above claim(s) \_\_\_\_\_ is/are withdrawn from consideration.
- 5) ☐ Claim(s) \_\_\_\_\_ is/are allowed.
- 6) ☒ Claim(s) 1,3-11,13,15,16 and 19-33 is/are rejected.
- 7) ☐ Claim(s) \_\_\_\_\_ is/are objected to.
- 8) ☐ Claim(s) \_\_\_\_\_ are subject to restriction and/or election requirement.

**Application Papers**

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on \_\_\_\_\_ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.  
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).  
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

**Priority under 35 U.S.C. § 119**

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All    b) ☐ Some    \* c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
  2. ☐ Certified copies of the priority documents have been received in Application No. \_\_\_\_\_.
  3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

\* See the attached detailed Office action for a list of the certified copies not received.

**Attachment(s)**

- |  |   |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892)                                | 4) <input type="checkbox"/> Interview Summary (PTO-413)<br>Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)                       | 5) <input type="checkbox"/> Notice of Informal Patent Application                       |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)<br>Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____  |

### **DETAILED ACTION**

1. Claims 1, 3-11, 13, 15, 16 and 19-33 are rejected and claims 2, 12, 14, 17 and 18 have been cancelled.

### ***Continued Examination Under 37 CFR 1.114***

2. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 2 August 2007 has been entered.

### ***Claim Rejections - 35 USC § 103***

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. **Claims 1, 3-11, 13, 15, 16 and 19-33 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No 7,206,836 to Dinker et al (hereafter Dinker) in view of the article "The Anatomy of a Large-Scale Hypertextual Web Search Engine" by Brin et al (hereafter Brin) in view of US Patent No 5,689,706 to Rao et al (hereafter Rao).**

**Referring to claim 1**, Dinker discloses a file system, comprising:

a plurality of servers [nodes 101A-101C] configured to store file data as chunks [subsets] (see column 3, lines 31-46 and Fig 1A); and

a master [replication topology manager] connected to the servers (see column 5, lines 47-59)

where the master is configured to:

communicate with the servers upon startup of the master to identify the chunks stored by the servers, and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location (see column 6, lines 8-67).

Dinker fails to explicitly disclose the further limitation wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks, store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Brin discloses search engine architecture for storing web pages, including the further limitation of wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks [doc ID], store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond [word ID], and store location data that identifies which of the servers [barrels] stores which of the chunks (see Section 4.1 – 4.2.7).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker to save the metadata disclosed by Brin. One would have been motivated do in order to track which files need to be copied from one node to another when there is a node failure.

The combination of Dinker and Brin (hereafter Dinker/Brin) fails to explicitly disclose the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Rao discloses managing a distributed system with replicated files, including the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data (see column 9, lines 8-25).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker/Brin to save the log disclosed by Rao. One would have been motivated do in order to provide accurate information pertaining to the location of files by tracking location changes.

**Referring to claim 3**, the combination of Dinker and Brin (hereafter Dinker/Brin/Rao) discloses the system of claim 1, wherein the master is further configured to control placement of new chunks at the servers (Dinker: see column 6, lines 8-49).

**Referring to claim 4**, Dinker/Brin/Rao discloses the system of claim 3, wherein when controlling the placement of new chunks, the master is configured to: identify one or more of the servers to store the new chunks based on at least one of utilization of the

Art Unit: 2167

servers, prior chunk distribution involving the servers, network topology, or failure correlation properties associated with the servers, and place the new chunks at the identified one or more servers (Dinker: see column 6, lines 8-67).

**Referring to claim 5**, Dinker/Brin/Rao discloses the system of claim 1, wherein the master is further configured to control redistribution of the chunks stored by the servers (Dinker: see column 6, lines 27-38).

**Referring to claim 6**, Dinker/Brin/Rao discloses the system of claim 5, wherein when controlling redistribution of the chunks, the master is configured to: select a chunk to redistribute based on a current distribution of the chunks, identify one or more of the servers to which to move the selected chunk, and move the selected chunk to the identified one or more servers (Dinker: see column 6, lines 8-67).

**Referring to claim 7**, Dinker/Brin/Rao discloses the system of claim 1, wherein the master is further configured to monitor a state [status] of the servers [nodes] (Dinker: see column 6, lines 13-16).

**Referring to claim 8**, Dinker/Brin/Rao discloses the system of claim 7, wherein the master is configured to exchange heartbeat signals with the servers to determine the state of the servers (Dinker: see column 6, lines 13-15).

**Referring to claim 9**, Dinker/Brin/Rao discloses the system of claim 8, wherein the heartbeat signals include space utilization information [how often each subset of data is accessed] (Dinker: see column 6, lines 40-44).

**Referring to claim 10**, Dinker/Brin/Rao discloses the system of claim 7, wherein the state of the servers includes information regarding the chunks stored by the servers (Dinker: see column 6, lines 8-67).

**Referring to claim 11**, Dinker/Brin/Rao discloses the system of claim 10, wherein the information includes version numbers of the chunks (Dinker: see column 6, lines 8-67).

**Referring to claim 13**, Dinker discloses a master in a file system that includes the master connected to a plurality of servers, the master comprising:

means for communicating with the servers to identify the file data stored by the servers as chunks, and means for storing in a non-persistent manner, location information that identifies ones of the servers that store the chunks (see column 6, lines 8-67).

Dinker fails to explicitly disclose the further limitation wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks, store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Brin discloses search engine architecture for storing web pages, including the further limitation of wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks [doc ID], store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond

[word ID], and store location data that identifies which of the servers [barrels] stores which of the chunks (see Section 4.1 – 4.2.7).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker to save the metadata disclosed by Brin. One would have been motivated do in order to track which files need to be copied from one node to another when there is a node failure.

The combination of Dinker and Brin (hereafter Dinker/Brin) fails to explicitly disclose the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Rao discloses managing a distributed system with replicated files, including the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data (see column 9, lines 8-25).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker/Brin to save the log disclosed by Rao. One would have been motivated do in order to provide accurate information pertaining to the location of files by tracking location changes.

**Referring to claim 15**, the claim is rejected on the same grounds as claim 1.

**Referring to claim 16**, the claim is rejected on the same grounds as claim 1.

**Referring to claim 19**, Dinker/Brin/Rao discloses the file system of claim 1, where the file identifiers are organized hierarchically in a tree of directories (Brin: see Section 4.1: Google Architecture Overview).



**Referring to claim 20**, Dinker/Brin/Rao discloses the file system of claim 1, where the master stores the namespace data using prefix-compression [compression] (Brin: see Section 4.2.2 Repository).

**Referring to claim 21**, Dinker/Brin/Rao discloses the file system of claim 1, where the master is configured to identify one of the chunks via a chunk handle that uniquely identifies the one of the chunks [word ID] (Brin: see Section 4.1: Google Architecture Overview).

**Referring to claim 22**, Dinker/Brin/Rao discloses the file system of claim 21, where the chunk handle encodes a timestamp (Brin: see Section 4.2.2 Repository).

**Referring to claim 23**, Dinker/Brin/Rao discloses the file system of claim 1, where the master is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers (Dinker: see column 6, lines 8-67).

**Referring to claim 24**, Dinker/Brin/Rao discloses the file system of claim 1, where the operation log includes a logical timeline that defines an order for concurrent operations (Rao: see column 9, lines 8-25).

**Referring to claim 25**, Dinker/Brin/Rao discloses the file system of claim 1, where the master is configured to: determine when a size of the operation log exceeds a threshold, and create a checkpoint of the operation log when the size of the operation log exceeds the threshold (Rao: see column 11, lines 11-34).

**Referring to claim 26**, Dinker/Brin/Rao discloses the file system of claim 25, where the master is configured to: create a new operation log file, and create the checkpoint as a background operation (Rao: see column 11, lines 11-34).

**Referring to claim 27**, Dinker/Brin/Rao discloses the master of claim 13, where the operation log includes a logical timeline that defines an order for concurrent operations (Rao: see column 9, lines 8-25).

**Referring to claim 28**, Dinker/Brin/Rao discloses the master of claim 13, further comprising: means for determining when a size of the operation log exceeds a threshold, and means for creating a checkpoint of the operation log when the size of the operation log exceeds the threshold (Rao: see column 11, lines 11-34).

**Referring to claim 29**, Dinker/Brin/Rao discloses the master of claim 28, where the means for creating the checkpoint includes: means for creating a new operation log file, and means for creating the checkpoint as a background operation (Rao: see column 11, lines 11-34).

**Referring to claim 30**, Dinker discloses a method performed by a master in a file system that includes the master device connected to a plurality of server devices, the method comprising:

communicating with the server devices to identify the file data stored by the server devices as chunks (see column 6, lines 8-67).

Dinker fails to explicitly disclose the further limitation wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks, store mapping data that maps the file identifiers to the chunks

to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks. Brin discloses search engine architecture for storing web pages, including the further limitation of wherein the master is configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks [doc ID], store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond [word ID], and store location data that identifies which of the servers [barrels] stores which of the chunks (see Section 4.1 – 4.2.7).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker to save the metadata disclosed by Brin. One would have been motivated do in order to track which files need to be copied from one node to another when there is a node failure.

The combination of Dinker and Brin (hereafter Dinker/Brin) fails to explicitly disclose the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Rao discloses managing a distributed system with replicated files, including the further limitation of the master configured to store an operation log that includes a record of changes to at least one of the namespace data or the mapping data (see column 9, lines 8-25).

It would have been obvious to one of ordinary skill in that art at the time of the invention to utilize the master of Dinker/Brin to save the log disclosed by Rao. One

would have been motivated do in order to provide accurate information pertaining to the location of files by tracking location changes.

**Referring to claim 31**, Dinker/Brin/Rao discloses the method of claim 30, where maintaining the operation log includes storing a logical timeline that defines an order for concurrent operations (Rao: see column 9, lines 8-25).

**Referring to claim 32**, Dinker/Brin/Rao discloses the method of claim 30, further comprising: determining when a size of the operation log exceeds a threshold, and creating a checkpoint of the operation log when the size of the operation log exceeds the threshold (Rao: see column 11, lines 11-34).

**Referring to claim 33**, Dinker/Brin/Rao discloses the method of claim 32, where creating the checkpoint includes: creating a new operation log file, and creating the checkpoint as a background operation (Rao: see column 11, lines 11-34).

**Conclusion**


Any inquiry concerning this communication or earlier communications from the examiner should be directed to Kimberly Lovel whose telephone number is (571) 272-2750. The examiner can normally be reached on 8:00 - 4:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Cottingham can be reached on (571) 272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Kimberly Lovel  
Examiner  
Art Unit 2167

26 October 2007  
kml

  
JOHN COTTINGHAM  
SUPERVISORY PATENT EXAMINER  
TECHNOLOGY CENTER 2100